

Generating Optimal Initial Conditions for Smooth Particle Hydrodynamics Simulations

S. Diehl^{1,2}, G. Rockefeller², C. L. Fryer^{*2}, D. Riethmiller³ & T. S. Statler^{3,4}

¹*Nuclear & Particle Physics, Astrophysics & Cosmology Group (T-2), Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, NM, 87545*

²*Computational Methods (CCS-2), Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, NM, 87545*

³*Astrophysical Institute, Ohio University, Athens, OH, 45701*

⁴*National Science Foundation, USA*

5 November 2012

ABSTRACT

We present a new optimal method to set up initial conditions for Smooth Particle Hydrodynamics (SPH) simulations, which may also be of interest for N-body simulations. This new method is based on weighted Voronoi tessellations (WVTs) and can meet arbitrarily complex spatial resolution requirements. We conduct a comprehensive review of existing SPH setup methods, and outline their advantages, limitations and drawbacks.

Key words: Smooth Particle Hydrodynamics – SPH – N-body – Simulations – Initial Conditions

1 INTRODUCTION

Smooth Particle Hydrodynamics (SPH) is a Lagrangian hydrodynamics modeling technique, that has been independently developed in the late 1970s by Lucy (1977) and Gingold & Monaghan (1977). In this grid-less technique, fluid elements are represented by individual particles, that act according to hydrodynamic flow equations. Subsequently, SPH has been used to model objects ranging from star formation (e.g. Monaghan & Lattanzio 1991; Bonnell & Bastien 1992; Whitworth 1998), planet formation (e.g. Benz et al. 1986; Mayer et al. 2002; Nelson & Benz 2003), cosmology (e.g. Navarro et al. 1995b,a; Springel 2005), stellar collisions (e.g. Benz & Hills 1987; Davies et al. 1991, 1992; Rasio & Shapiro 1991), stellar mergers (e.g. Davies et al. 1994; Fryer

* E-mail: fryer@lanl.gov

& Heger 2005; Lee et al. 2001; Rasio & Shapiro 1992; Terman et al. 1994; Rosswog et al. 1999; Yoon et al. 2007; Diehl et al. 2008; Motl et al. 2008), gas dynamics in the Galactic center (e.g. Rockefeller et al. 2005, 2004; Cuadra et al. 2005), galaxy mergers (e.g. Cox et al. 2006; Khalatyan et al. 2008; Thakar & Ryden 1998) to supernovae (Fryer & Warren 2002; Fryer et al. 2007; Hungerford et al. 2003), just to name a few examples.

Each SPH particle i has an associated size, its so-called smoothing length, which is usually denominated as h_i . Fluid properties such as temperature or density are smoothed according to a smoothing function W , which is referred to as the SPH kernel. The most commonly used kernel functions are cubic splines (Monaghan 1992) that are compact within 2 smoothing lengths, and are vanishing outside. A particle's fluid properties can then be expressed as a linear combination of all neighboring particles contributing to a particle's position. Thus, it is essential for this technique to start out with initial conditions whose interpolation properties are as accurate as possible. In additions, the initial particle setup has to be as natural to SPH as possible, i.e. it should be very close to a configuration that would arise by itself in an SPH simulation. This also implies that we need a particle setup that is not based on lattice structures, as the artificial alignments of SPH particles may cause numerical artifacts.

The simplest and also most popular, despite their known drawbacks, setup methods include the arrangement in lattices due to their well known interpolation accuracies and ease of setup. The simplest and one of the most popular methods is the cubic lattice configuration, which has been shown to be a non-stable equilibrium configuration and has strong preferred geometrical axes along the x, y and z axes (Lombardi et al. 1999). There are infinitely many lattice configuration that could in principle be used to produce SPH initial conditions. During the rest of this paper, we will focus on the two arguably best lattice configurations, cubic and hexagonal close packing. These two optimal and most efficient ways to pack spheres of equal sizes are stable against random perturbation and thus much preferred to a simple cubic lattice (Monaghan 1992). We will also include a comparison with a new configuration method based on a quaquaversal tiling of space that has recently been suggested for quasi-random initial conditions of cosmological N-body simulations (Hansen et al. 2007).

To avoid geometrical effects, initial conditions often get slightly perturbed and then relaxed into a stable configuration by applying a pseudo dampening force that is proportional to but directed against the particle velocities (e.g. Rosswog et al. 2008). While this additional step before calculation is perfectly acceptable to produce low-noise initial conditions, it is computationally rather expensive and usually only applicable for initial conditions that are static in nature, as the

net forces on the whole set of particles should vanish. In addition, there is no way to guarantee the exact configuration into which the particles will settle at the end.

Another major problem in setting up initial conditions for SPH is that many astrophysical situations require very large dynamic ranges in density, for example in modeling stars or galaxies. These problems tend to require large dynamic ranges in resolution as well, as large ranges of particle masses in SPH are undesirable. Yet, setting up arbitrary initial conditions with a spatially varying resolution is an unsolved problem so far, with very few proposed solutions so far that are only applicable in spherical symmetry (e.g. Rosswog et al. 2008; Fryer et al. 2007).

In this paper, we will propose a solution to this problem based on weighted Voronoi tessellations (WVTs), and present a new method to set up SPH initial conditions with arbitrary, spatially varying resolution requirements in 2, 3 or N dimensions. In short, we will refer to this new method as the “WVT setup.” We will start out with laying out the requirements of an optimal setup technique in section 2 and then go on to summarize and compare existing popular particle setup methods (§3). To the best of our knowledge, this is the first comprehensive comparison of SPH setup techniques, despite the known importance of initial conditions for SPH simulations. We then introduce our new WVT technique in section 4, and demonstrate its capabilities on various different examples in section 5. In §6 we will quantitatively compare this new setup method with existing technique, before we conclude in section 7.

2 REQUIREMENTS FOR AN OPTIMAL PARTICLE CONFIGURATION

First, let us define some obvious key requirements that an optimal SPH particle configuration setup has to fulfill:

Isotropy. The particle configuration should be locally and globally isotropic, i.e. not impose any particular preferred direction at any location of the simulation domain. The main reason for this requirement is the fact that shocks moving along a perfectly aligned string of SPH particles behave differently than in other directions (Herant 1994). In addition, spatially correlated density perturbations can excite modes along these preferred directions.

High Interpolation Accuracy. The setup should be locally uniform to minimize noise in the density interpolation. Ideally, for a uniform resolution, the interpolation accuracy should be comparable to that of perfectly uniform lattice configurations. This interpolation accuracy should also worsen for non-uniform particle configurations. Any deviations should also be isotropic and have

no preferred directions, in order not to excite non-physical modes in the simulation domain. This requirement is equivalent to enforcing a low particle noise.

Versatility. The ideal method should be able to reproduce any spatial configuration, and not pose any constraints on a particular symmetry. In particular, this requires the setup to work with interpolation of a tabulated data set and not require analytical solutions.

Ease of Use. Ideally, the algorithm should either be publicly available as a stand-alone routine or be easy to implement on top of any existing SPH code.

3 POPULAR PARTICLE SETUP METHODS

Since the invention of the SPH technique, many different methods have been employed to set up initial conditions in multiple dimensions. In this section, we will summarize all particle setup methods known to us, or that have been described in the literature. Readers that are already familiar with these methods, or are just interested in details of the new WVT setup method, may skip this section entirely or choose to take a look at Figure 1 for a simple comparison of a sphere with approximately 22,000 particles targeted at being uniformly distributed. Figure 2 shows an equivalent figure for spatially adaptive configurations.

3.1 Spatially Uniform Distributions

We first focus on setup methods that have been employed to set up uniformly distributed particle configurations.

Cubic Lattice (CL). Probably by far the simplest and fastest way to set up a uniform particle distribution is to arrange them on a cubic lattice. This method has received widespread use early-on in both SPH (Monaghan 1992) and N-body simulations (Efsthathiou et al. 1985). One of the obvious problems with this method is that it has very pronounced preferred directions along the x, y, and z-axis, and their diagonals, as can be easily seen in the upper left corner of Figure 1. In addition, the cubic lattice structure is not a stable equilibrium configuration when the particles are perturbed (Lombardi et al. 1999), as there are other more compact particle configurations that are energetically favorable, such as cubic or hexagonal close packing schemes.

Cubic Close Packing (CCP). A more compact lattice structure is produced when one inserts an additional particle into the center of each of the 6 faces of the cubes in the cubic lattice. This

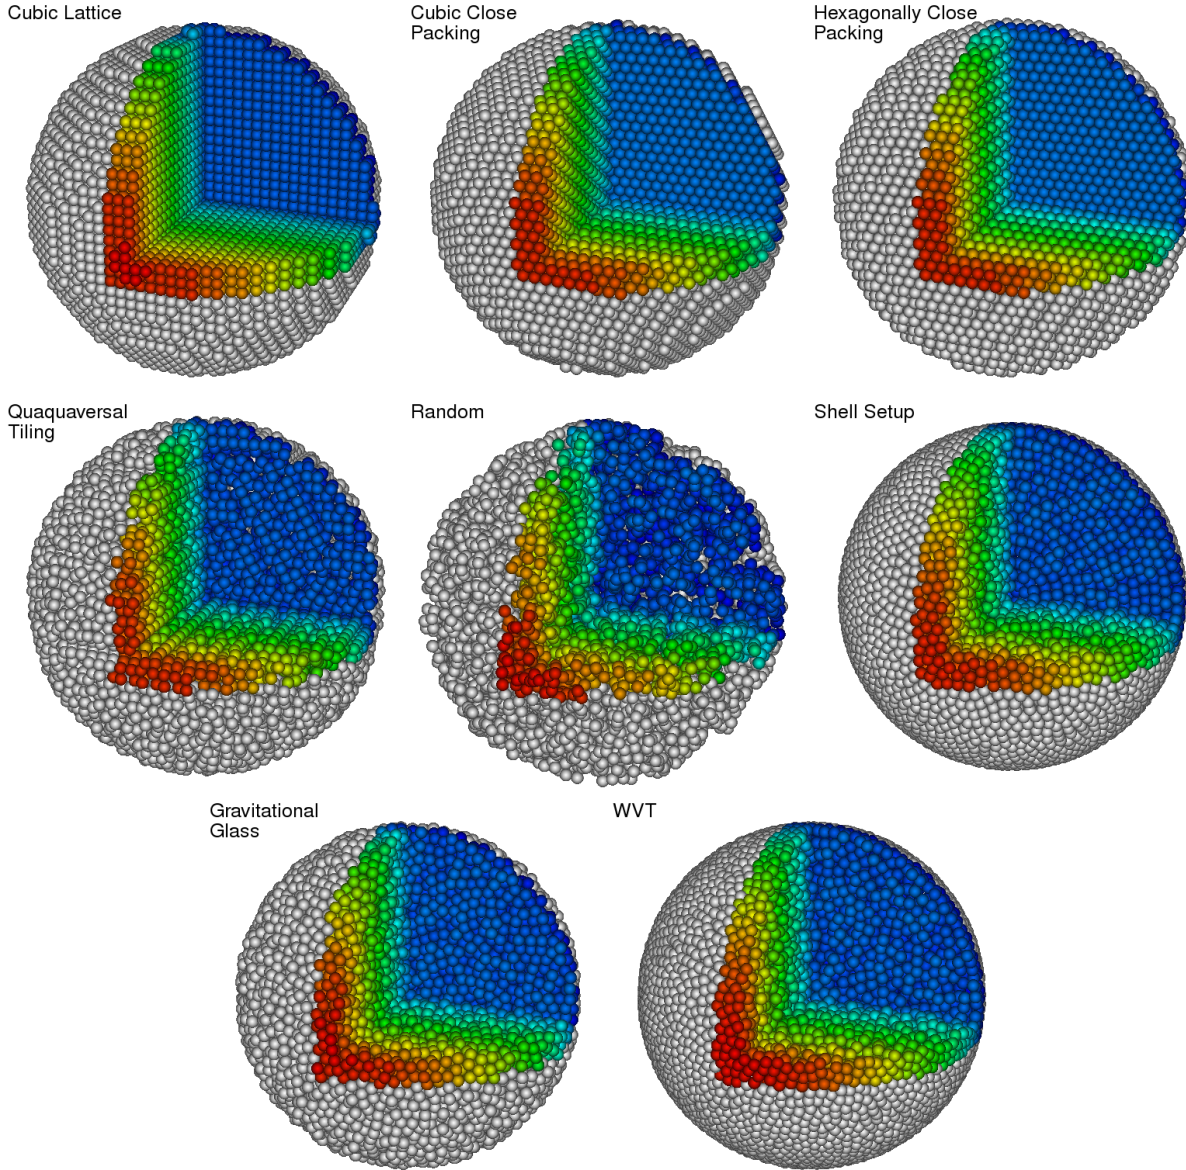


Figure 1. Popular configurations for setting up spatially uniform SPH initial conditions. From the top left corner to the bottom right: cubic lattice, cubic close packing, hexagonal close packing, quaquaversal tiling, random configuration, concentric shell setup, gravitational glass, and the new WVT approach. All images show approximately the same number of particles in the sphere (22,000). One quadrant of the sphere is cut out to allow a view into the inner configuration. Colors change along the z-axis to produce a better 3-d depth.

results in the well-studied cubic close packing (CCP) configuration, also known as face-centered cubic close packing (Figure 1, top center panel). This configuration is one of the optimal ways to pack uniform spheres together, with a packing density of 74%. Similar to the cubic lattice, it has the problem of having preferred directions along the axis and diagonals of the lattice, and multiple other planes in which particles are arranged in a hexagonal grid. However these preferred direction are much less pronounced than for the cubic lattice configuration. The simplest way to construct this configuration is to start with a plane with spheres in a hexagonal configuration (plane A), and lay on top another such plane (B) so that the spheres fit into the suppressions created by the layer

of spheres from the lower plane, filling half the suppressions. The third such plane (C) will then be oriented in a way to fill the other half of the suppressions of plane A, while at the same time fitting into the suppressions of plane B. This pattern is then continuously repeated to produce an ABCABC order.

Hexagonal Close Packed Lattice (HCP). A very similar lattice configuration is the second optimal packing scheme of uniform spheres, hexagonal close packing (HCP), as seen in the top right panel of Figure 1. HCP is equally dense and optimal as CCP, with very similar properties. The only difference in its construction is that instead of the ABC pattern as in the CCP lattice, the every second layer of hexagonal lattice planes is identical, resulting in an ABAB pattern. Due to their relatively simple implementation, close packing schemes have been utilized in many different SPH settings (e.g. Davies et al. 1991, 1992).

Quaquaversal Tiling (QVT). Another lattice-like particle configuration has been introduced by Hansen et al. (2007)¹ based on a quaquaversal tiling (QVT) of space (Conway & Radin 1998). QVT hierarchically tiles the 3d space into triangular prisms that are rotated about orthogonal axes $2/3\pi$ and $1/2\pi$. Originally introduced as a way to set up cosmological initial condition, recent work by Wang & White (2007) argues against this choice. As can be seen in the middle left panel of Figure 1, this setup has many characteristics of a grid.

Gravitational Glass (GG). The best way to set up cosmological initial conditions however is the generation of a gravitational glass Wang & White (2007). This method simply reverses the sign of gravity and lets the particles settle into an equilibrium configuration while dampening their motion. In this paper, we use the implementation of a gravitational glass as provided in the publicly-available Gadget2 code. This method is particularly effective when used with periodic boundary conditions. A cube with a fixed number of particles can then be replicated numerous times to achieve a larger total number of particles in the final configuration. While this is disadvantageous for cosmological N-body simulations, which require a certain degree of homogeneity on all scales, SPH only requires locally optimal configurations and is not affected by this. Thus, a single instance of a gravitational glass can be concatenated to effectively produce arbitrarily large particle configurations.

¹ publicly available at
<http://www-theorie.physik.unizh.ch/~hansen/qua>

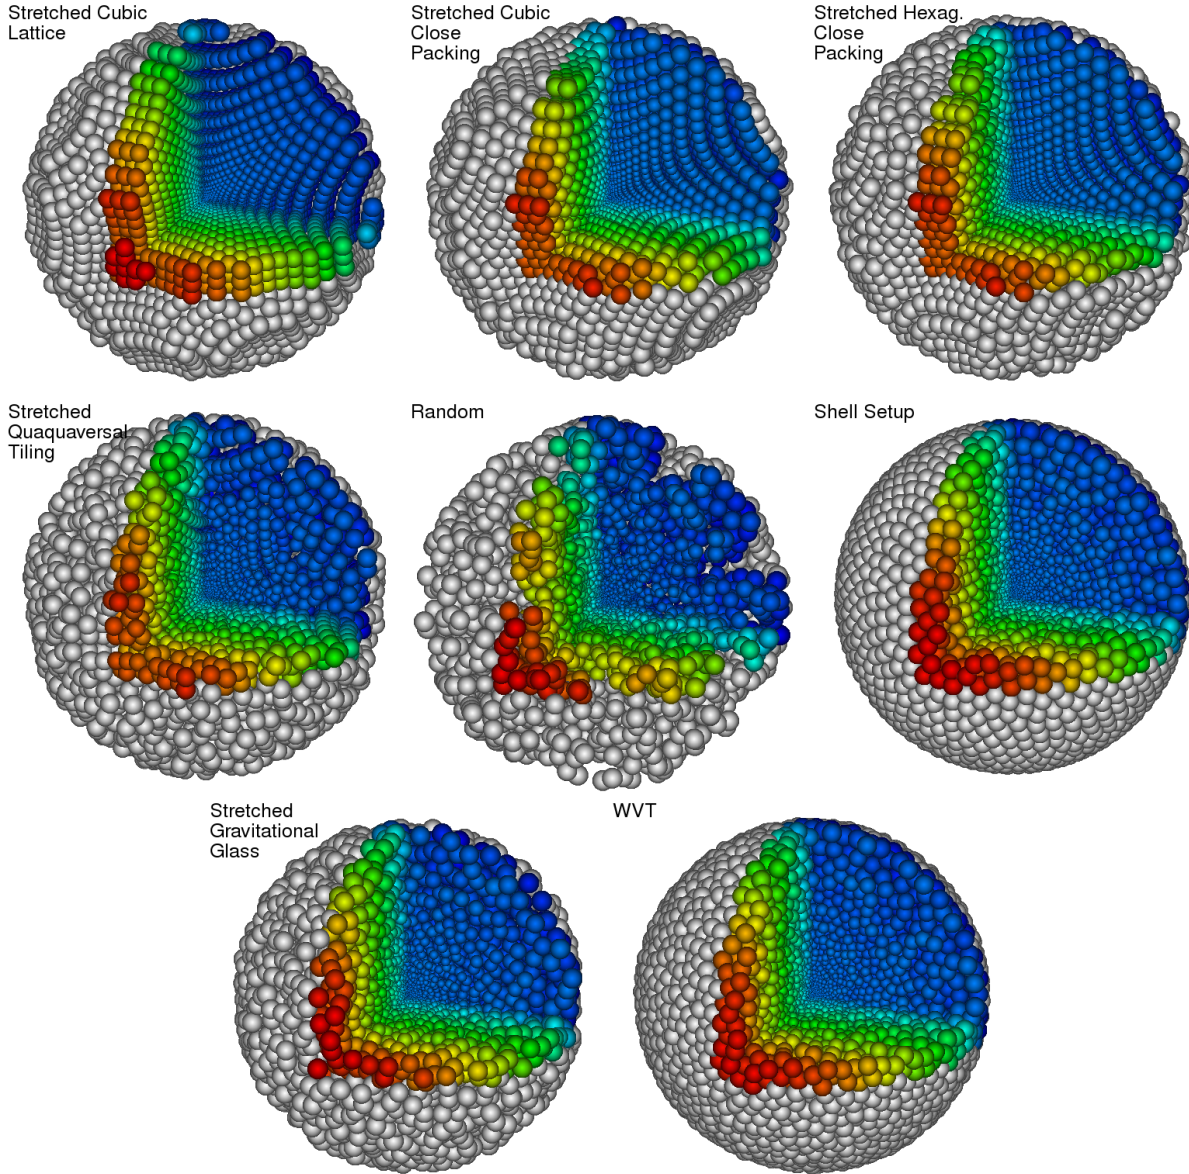


Figure 2. Popular configurations for setting up spatially adaptive SPH initial conditions. From the top left corner to the bottom right: stretched cubic lattice, stretched cubic close packing, stretched hexagonal close packing, stretched quaquaversal tiling, random configuration, concentric shell setup, stretched gravitational glass, and the new WVT approach. All images show approximately the same number of particles in the sphere (22,000), and the particles' sizes reflect the desired particle spacing. One quadrant of the sphere is cut out to allow a view into the inner configuration. Colors change along the z-axis to produce a better 3-d depth.

3.2 Spatially Adaptive Distributions

We now discuss setup methods that are capable of producing spatially adaptive particle distributions. To the best of our knowledge, all of these methods have so far been employed to create spherically symmetric configurations. An example for each of these methods is shown in Figure 2 for a direct qualitative comparison.

Random Configuration (RC). The simplest possibility to produce a spatially adaptive particle configuration is to distribute particles randomly according to an underlying probability distribution. For example, Terman et al. (1994) have used this technique in combination with the relaxation method to create adaptive initial conditions. This is absolutely necessary, as in SPH, this method results in very clumpy distributions with very low interpolation accuracy. We only mention this method as it represents the 0th order distribution, and because it is the starting point for our WVT iterations as we describe in detail in section 4.

Stretched Lattice (SL). To achieve a spatially adaptive resolution, Herant (1994) and later Rosswog et al. (2008) proposed to stretch a uniform lattice configuration in the radial direction. With this method, each point coordinate \mathbf{r} of the uniform lattice is multiplied by a radially varying scaling factor $q(r)$ to achieve the desired spherically symmetric distribution, such that $\mathbf{r}' = q(r) \mathbf{r}$. This also implies through simple geometry that the given distance δ between two particle on the shell with radius r is now also scaled by $q(r)$, effectively setting the resolution to $\delta'(r') = q(r) \delta$.

Thus, the problem has now been reduced to figuring out how to choose $q(r)$ to produce the desired resolution in the new stretched coordinates, i.e. $\delta'(r')$ has to obey the differential equation $r' - \delta'(r') r = 0$. While it is entirely possible to solve this problem analytically for simple δ' functions by substituting r' and solving, it is more convenient to use a more generally applicable technique to find the root of the function $f(r') = r' - \delta'(r') r$, with its derivative derivative $df/dr' = 1 - d\delta'/dr' r$. We chose the simple Newton-Raphson technique by iterating over $r'_{n+1} = r'_n + f(r'_n) [df/dr'(r'_n)]^{-1}$. Note that r is a constant parameter in this context, as it is given by the known position of the particle in the uniform lattice, contrary to r' .

As this distorted lattice essentially incorporates and even aggravates all of the undesirable characteristics of a lattice configuration, it is essential *not* to use this stretched lattice configuration directly, but rather to relax the resulting configuration before using it in an actual SPH simulations (Rosswog et al. 2008).

Stretched Glass (SG). Instead of radially stretching or compressing a uniform lattice configuration, it is also in principle possible to generate a gravitational glass of uniform resolution and stretch this glass accordingly, to avoid the strong preferred lattice symmetry axes. To the best of our knowledge, this method has never been employed so far.

Concentric Shell Setup (CSS). In many supernova calculations using SNSPH (e.g. Fryer & Warren 2002; Fryer & Young 2007; Fryer et al. 2006, 2007; Hungerford et al. 2003), the initial

conditions are set using shell templates. For a given particle count, such a shell template is created by first randomly placing the particles in a shell of unit radius. The particles are given a repulsive force and then the entire system is evolved until the deviation in the particle separation falls below some tolerance, essentially creating a two-dimensional gravitational glass wrapped around a sphere. The templates are then used to match a given spherical density profile and either a resolution (or particle mass) requirement. The spherical object is constructed from the inside outward, each concentric shell determining the new position of the next shell (one smoothing length above the previous shell). The shells are randomly rotated and placed on top of each other so that, even if the same template is used, the setup is random. A variant of this method had originally been proposed by Herant (1994), but to our knowledge has never been extensively described in the literature before.

The advantage of this technique is that the particles are placed randomly and hence have no preferred axis. Depending on the tolerance set for the template creation, the resolution used, and the density gradient, this technique can match a spherical density profile to arbitrary precision. For low resolution core-collapse calculations, Fryer & Young (2007) achieved density perturbations below 3-5% (convection in the stellar models they were mapping argued for higher perturbations). For a high resolution mapping of an exploding star, Fryer et al. (2007) limited this perturbation to below 1%. This technique is tuned to spherical objects, and does not work, without major revision, on aspherical objects.

4 A NEW APPROACH USING 3-D WEIGHTED VORONOI TESSELATIONS

In the following section, we introduce a novel technique to generate spatially adaptive initial conditions for SPH simulations. In principle this technique is applicable in arbitrary dimensions, but for the sake of clarity, we will restrict our discussion to two and three dimensions. A Voronoi diagram (or tessellation) is a specific decomposition of a metric space defined by a set of generators. A cell in this space is defined by all particles closer (where closer is defined by a given metric) to a given generator than any other. In this way, the generators tessellate the space.

4.1 Weighted Voronoi Tessellations

Our new adaptive setup technique constructs approximate multiplicatively weighted Voronoi tessellations (WVTs, see Møller 1994, for more information), which do not impose any restrictions on the geometry of the desired configuration. A WVT is a tiling of space and is fully defined by

the locations \mathbf{z}_i of its set of N generators, and their associated scale lengths² δ_i . The boundary surface b between two bins i and j is then defined such that

$$|\mathbf{b} - \mathbf{z}_i|/\delta_i = |\mathbf{b} - \mathbf{z}_j|/\delta_j, \quad (1)$$

i.e. a piece of space is considered part of the bin to which its scaled distance is shortest. One can think of δ_i as factors that stretch or compress the metric inside the bins they belong to. For the special case that all scale lengths are equal, the space is tiled such that a bin encompasses all points that are closest to its generator, and the WVT is reduced to a normal Voronoi tessellation.

When the generators also coincide with the centroids of a WVT, the bins in the tessellation generally become very regular and round, and produce a so-called centroidal WVT. Placing SPH particles at these centroids produces the optimal configuration that we are striving for, as shown in the final panel of Figures 3 and 4. As there is no a priori analytical solution to achieve such an optimal, self-consistent distribution, this setup generally has to be derived iteratively. For centroidal Voronoi tessellations, this procedure is often referred to as the Lloyd algorithm (Lloyd 1982).

Our WVT setup method has been inspired by a two-dimensional adaptive binning technique recently developed by Diehl & Statler (2006), which in turn is based on previous work by Cappellari & Copin (2003).

4.2 Technique

One way to approach a centroidal WVT configuration is to start with with an initial particle configuration and iteratively converge to a solution. For the following discussion, let us assume that we are given a desired resolution $\delta(\mathbf{r})$ as a function of a three-dimensional position \mathbf{r} . In this context, “resolution” is referred to as a the desired average distance to a neighboring particle. The actual smoothing lengths h are then simply related by a constant proportionality factor that depends on the targeted number of neighbors N_{neigh} in the calculation, such that $h(\mathbf{r}) = N_{\text{neigh}}^{1/n} (\delta(\mathbf{r})/2)$ in n dimensions.

The most obvious and flexible choice to start out with is an initial random configuration according to the underlying particle probability distribution $P(\mathbf{r}) \propto h(\mathbf{r})^{-3}dV$ for a volume dV . We then evolve this configuration for multiple iteration steps by applying repulsive forces³ between the particles so they settle in the desired places. Figure 3 shows an instructive two-dimensional

² In principle, this scale length is allowed to be an asymmetric tensor; our discussion uses scalar weights for simplicity.

³ We implement forces between particles to iterate a solution with a given metric. Here, “force” is used as in enforcing a metric and should not be associated with a physical force in nature.

example for the whole iteration process for a uniform distribution of particles, starting with the probability distribution on the lower left and ending with the final product on the lower right. Figure 4 shows an equivalent sequence for an azimuthally symmetric, but non-uniform distribution.

To achieve the final configuration, these artificial forces have to fulfill three main properties.

(i) WVT bins only know about their direct neighbors, requiring these forces to vanish beyond a given radius. In practice, the bins should know about all the neighbors that they will feel in the actual SPH simulations to make the distribution and thus the interpolation properties smooth on that scale. This requires that the force is compact with 2 smoothing lengths h . In addition close and direct neighbors should have a stronger influence than second or third degree neighbors.

(ii) All “forces” on any particle should cancel out to a given precision in the desired equilibrium configuration, i.e when the distance between two particles is identical (or very close) to the desired resolution.

(iii) The “forces” should be such that the net displacement from the last position is on the order of a fraction of the desired resolution. This will ensure equally fast convergence for high and low-resolution portions of the setup. We will thus restrict our discussion from now on to the term “displacements” directly rather than talking about pseudo “forces”, as they do not have a physical meaning anyway.

Thus, equation (1) dictates the value of the repulsive “displacement” from bin j on bin i , which we express for simplicity directly as a net displacement

$$\Delta \mathbf{x}_i = \sum [\mu h_i f(h_{ij}, r_{ij}) \hat{\mathbf{r}}_{ij}] = \mu h_i \sum [f(h_{ij}, r_{ij}) \hat{\mathbf{r}}_{ij}], \quad (2)$$

with $h_{ij} = (h_i + h_j)/2$. The function $f(h_{ij}, r_{ij})$ should be compact within $2h$, and empirical tests show fastest convergence of the method for a r^{-2} dependence. In practice, we add an ϵ term in the denominator to avoid numerical problems for close particles, and subtract a constant value to make sure the function value vanishes at the boundary at $2h$ and is set to 0 if the separation is larger than that. Thus, we can express the function as $f(h_{ij}/r_{ij}) = [h_{ij}/(r_{ij} + \epsilon)]^2 + \text{const}$. The value of μ in equation (2) regulates what fraction of h_{ij} the particles are allowed to move during each iteration step. This free parameter should be chosen to ensure fast convergence. In practice, we shrink μ monotonically with the number of iterations, so particles can move relatively freely at the beginning and at the end “freeze” into their final position. Note that the particle spacing $\delta(\mathbf{r})$ can be an almost arbitrary function of space, as long as its value does not change significantly across one particle spacing to ensure convergence of the method, i.e. $\Delta\delta(\mathbf{r})/\delta(\mathbf{r}) \ll 1$. If the method

converges properly, the particles positions are coincident with the generators of the multiplicatively weighted Voronoi centroidal tessellation with the smoothing length as multiplication factor.

Note that we have narrowed down to a very specific subclass of weighted Voronoi tessellations. We use a multiplicatively weighted tessellation (one could also use additively weighted functions). We have also chosen a specific function for our multiplicative factor. We also emphasize that the functional form of f should not be unique to solve this problem. We chose the r^{-2} dependence as it reproduces locally many desirable properties of a gravitational glass. A functional form based on the SPH kernel would be another natural choice for this problem, though we did not thoroughly test this possibility. Certainly a broad set of WVT solutions are available and many will not set up conditions ideal for SPH modeling.

This technique is different than the Lloyd algorithm (Lloyd 1982) solutions which calculate Voronoi Tessellations and reset the particle positions at the centroids, iterating to convergence. Such techniques can guarantee the co-location of the centroid and particle to the chosen tolerance. With a given metric, one can redistribute the particles to follow the desired resolution. Our method allowed us to easily prescribe the desired resolution, but it does not guarantee co-location of the particle and centroid. However, for our models, we find that the particle position and centroid are co-located within -0.3% of the smoothing length.

4.3 Practical Implementation

Here we provide readers practical advice on how to implement the WVT setup on top of their own existing SPH code in this short section. You may also contact the authors if you would like to have a large-scale run set up for your code, or have questions on how to add WVT capabilities to your own parallel SPH code.

How to normalize $h(\mathbf{r})$. In our description of the WVT method above, we have assumed that we know a priori the desired particle spacing $\delta(\mathbf{r})$ as a function of position. However, for an arbitrary geometrical configuration, it is often difficult to know the required particle spacings to achieve a desired number of SPH particles, which is a more desirable quantity to specify very often. Thus, our implementation interprets the input particle spacing as a relative rather than an absolute value, and scales it according to a desired number of particles N_{SPH} and neighbors N_{neigh} . At each iteration step, $h(\mathbf{r})$ gets evaluated for all particle positions and we then compute the sum of all individual SPH particle volumes: $V_{\text{SPH}} = \sum_i [(4\pi/3)(2h_i)^3]$. Since we do know the actual

volume V of our computational domain and that we desire N_{neigh} neighbors within $2h$ for each particle, we scale all $h(\mathbf{r})$ values accordingly such that $V = V_{\text{SPH}}/N_{\text{neigh}}$.

How to treat boundaries. Most modern SPH codes have some kind of boundary treatment already implemented. Fixed boundaries are usually implemented by means of ghost particles (e.g. Herant 1994) that exert antisymmetric forces on the particles to keep them across a given boundary. WVT works well with this type of boundary treatment, and we suggest mirroring SPH particle layers within 2 smoothing lengths at the boundary interface and adding them to the set of normal particles during the pseudo force calculation step. Periodic boundaries may also be used, and the region of interest can simply be “cut out” afterwards. We find this method to work well for arbitrary geometries.

How to update particle positions. The most convenient way to implement WVT is to use the entire structure of your existing SPH code with as few changes as possible to the code. We suggest to modify the existing SPH loop to calculate the sum in equation (2) and then multiply this pseudo “velocity” by the “individual time steps” μh_i to get the particle updates. If the distribution does not appear to converge, we suggest decreasing the value of μ with each iteration.

When to stop the iterations. It is rather subjective to judge when an initial condition setup can be considered sufficiently good, and this decision is strongly dependent on the type of application. In our experience, slowly reducing the value of μ (the maximum fractional distance a particle can be moved in one time step, see equation 2) works extremely well and, for the cases we’ve studied (Raskin et al. 2009, 2010; Fryer et al. 2010; Passy et al. 2011; Ellinger et al. 2011) which include setups using from 100,000 to 50 million particles, convergence occurs in about 100 iterations, usually even after only 40 iterations. However, this number may strongly depend on the details of the algorithm, and the desired interpolation accuracy, an issue that we discuss in more detail in section 6.

5 EXAMPLE APPLICATIONS

Different applications may pose very different requirements on the resolution of a particular object. WVT has been designed to deal with exactly this issue. Let us now consider different examples, mostly from stellar interactions, that pose very different numerical requirements.

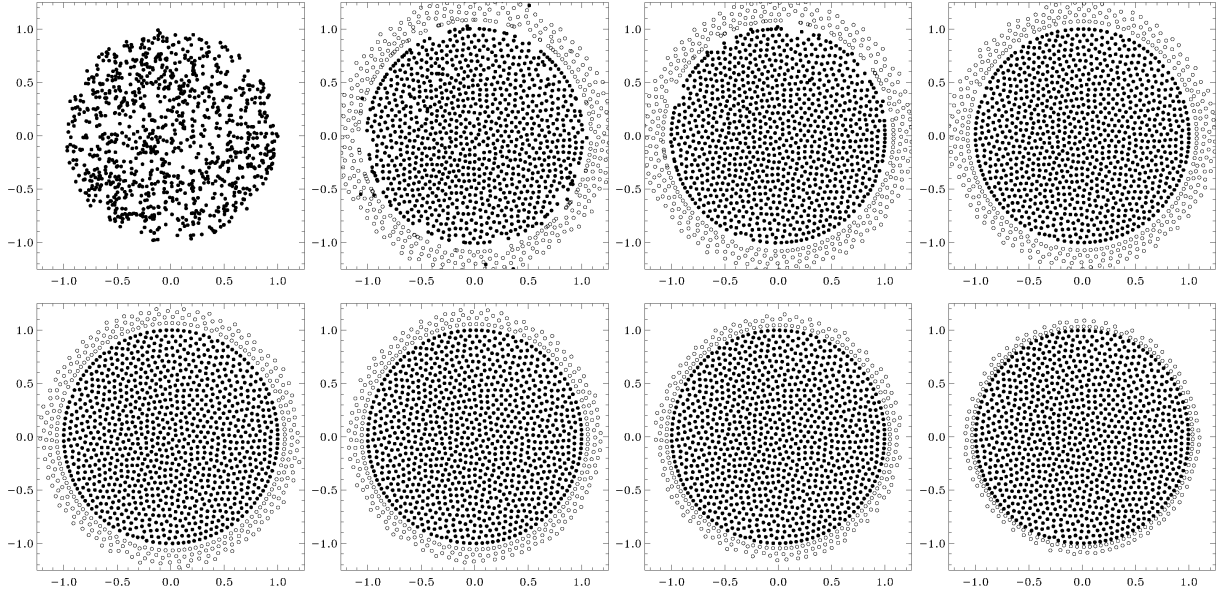


Figure 3. Two-dimensional example for producing a uniform particle density in a circle with radius 1 with the new WVT setup for 1000 particles. The frames show snapshots of the WVT iterations, starting with random initial conditions (top left), and then showing every 10th iteration, until the final setup is reached in the lower right panel (here, iteration 70). Note how the particles quickly converge toward a stable configuration, with “defects” slowly disappearing until a very uniform distribution is reached, and the outer layers describe a perfect circle. The hollow particles are “ghost particles” that are needed to ensure proper boundaries.

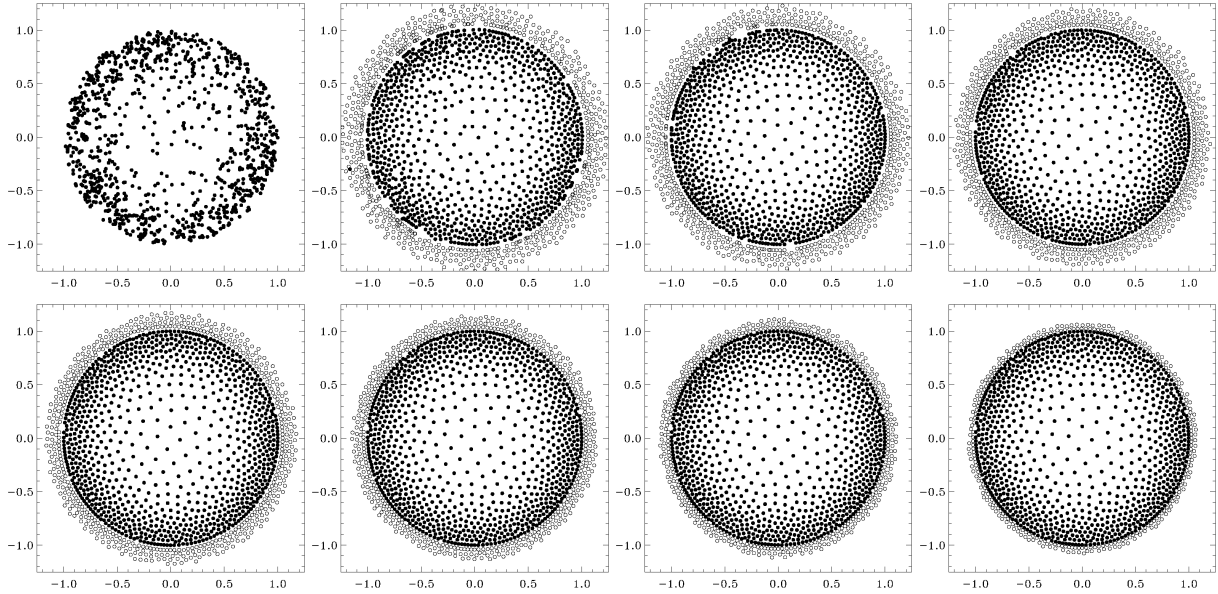


Figure 4. Same as Figure 3, but for a non-uniform particle distribution. The goal is to reach a 4 times higher particle resolution in the outer layers of the circle than in the center.

5.1 Uniform Particle density

In the first application, we consider representing a star with a uniform particle density, which would be appropriate for simple head-on collisions between two stars. In this situation, it is very likely that both center and outer layers will be heavily involved in the process, as all parts should

be equally affected during the merger, and would require equal resolution to resolve the hydrodynamical effects throughout the star.

Such a 2D particle setup is shown in the top left panel of Figure 5. The lower panel shows how well the desired resolution is achieved by measuring the average distance between particles for the closest 8 (red), 16 (green), 32 (blue) and 64 (orange) neighbors. The black solid line shows the expected theoretical values for 16 neighbors. The deviations at the boundary are due to the lack of neighbors across the boundary, and can of course be fixed by increasing h accordingly in that region.

5.2 More Resolution in the Center

However, if the user is interested in any type of mixing within the stars or during a merger, it is imperative to use equal-mass particles. Work by Lombardi et al. (1999) suggests that artificial forces between non-equal mass particles lead to numerical diffusion and artificial mixing. With the WVT setup, we can enforce equal mass particles by adjusting the particle separation according to the underlying density $\rho(\mathbf{r})$, such that $\delta(\mathbf{r}) \propto \rho(\mathbf{r})^{-1/3}$. This results in a setup with more resolution in the center, as shown in the center panels of Figure 5.

5.3 More Resolution in the Outer Layers.

If one is interested in studying the more gentle Roche Lobe overflow phase in a binary, one needs as much numerical resolution as possible in the outer layers of the donor star, in order to sufficiently resolve the overflow and accretion stream. The middle panel of Figure 5 shows a polytrope where we have put more resolution in the outermost layer than in the center.

5.4 Asymmetric Initial Conditions: Double Degenerate Binary

Figure 6 shows an example for an asymmetric, three-dimensional setup with WVT. The picture depicts a double degenerate binary system with the donor (right) on the verge of overflowing its Roche lobe. Note how the size of the SPH particles (varying sizes and colors of spheres) is much smaller in the outer layers of the donor, which helps significantly in resolving the mass transfer stream in the simulation (Motl et al. 2008). The evolution of the system is extremely sensitive to the initial mass transfer (as it governs the evolution of the orbit). Resolving this mass transfer is critical to achieving good agreement between the SPH and rotating grid simulations.

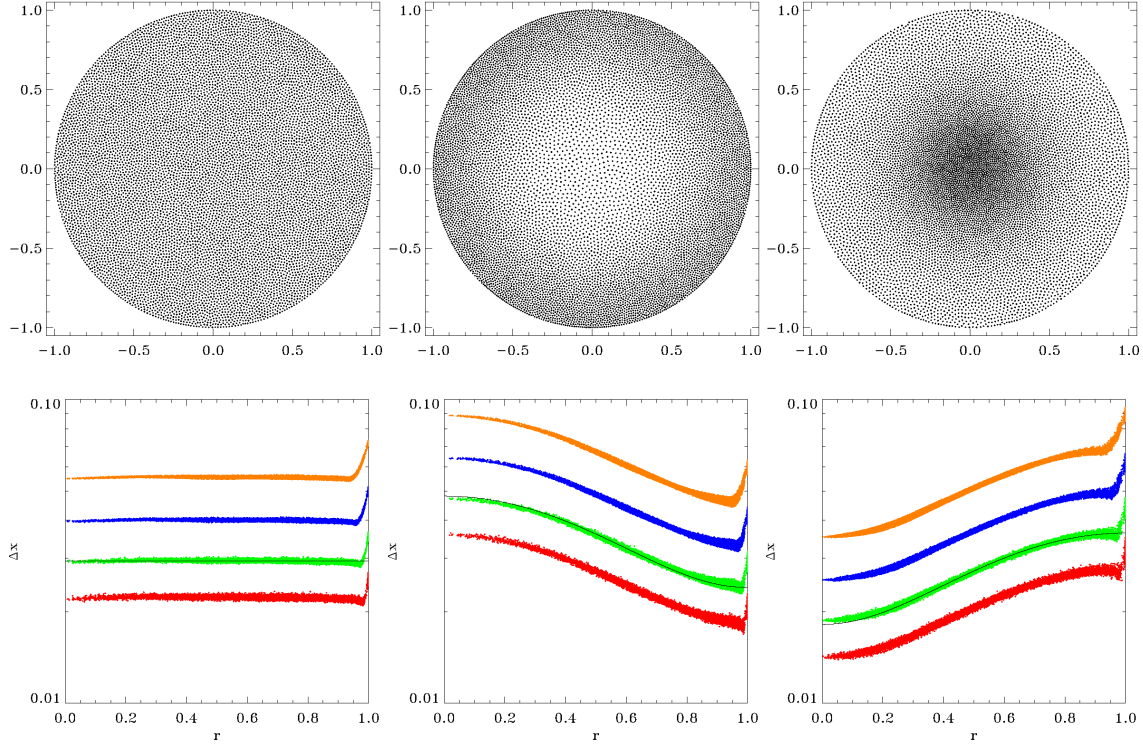


Figure 5. Top panels: particle configurations in two-dimensional examples. We consider three different configurations: uniform particle density (left), more resolution in outer layers (center), more resolution at center (right). The bottom panels show the actual particle separations as a function of radius. The points measure the average distance to the closest 8 (red), 16 (green), 32 (blue) and 64 (orange) neighbors. The solid black line shows the input resolution scaled for the closest 16 neighbors, indeed closely following the green data points.

5.5 Asymmetric Initial Conditions: Elliptical Galaxies

Simulations of normal elliptical galaxies commonly use a gas component embedded within a dark matter halo in order to model evolution properly. Creating an initial configuration for a simulation of feedback effects in such a galaxy, it is reasonable to say the gas has settled hydrostatically into the dark matter potential and is quiescent. We assume that the gravitational effects of the gas are small compared to those of the dark matter. Then, we should expect the gas density ρ_{gas} to follow the dark matter potential Φ_{DM} , assuming a polytropic equation of state:

$$\rho_{\text{gas}}(x, y, z) = \left[\frac{-1}{K(n+1)} \Phi_{\text{DM}}(x, y, z) + C \right]^n \quad (3)$$

where K is a constant, n is the polytropic index, and C is a constant of integration that determines the sharpness of the edge of the density distribution. Note that ρ_{gas} and Φ_{DM} need not be spherical in shape; the dark matter potential maps into the gas density regardless of its degree of eccentricity.

In fact, we may take direct advantage of the non-requirement for spherical shape; the WVT code contains an option to use a cloud of dark matter particles as 3-dimensional interpolation points in order to determine the value of the dark matter potential at a given position. Then, Equation 3 defines the mass density for a gas particle placed at that point.

If Equation 3 is applied correctly, the resulting surfaces of constant gas mass density must coincide with the surfaces of constant dark matter potential. Thus, as a diagnostic test, we first construct a self-consistent Hernquist sphere of N-body particles with a distribution function of Ossipkov-Merritt form (Osipkov 1979; Merritt 1985a,b; Kazantzidis et al. 2004) and anisotropy radius $r_a = 1 \times 10^{10}$, and then use the method of Holley-Bockelmann et al. 2001 and Widrow 2008 to deform the sphere adiabatically into a triaxial system. The resulting configuration of N-body particles has approximate axis ratio 17:15:14, and forms the set of interpolation points from which to create a cloud of gas particles via WVT. The example gas cloud analyzed in Figure 7 has a polytropic index of $n = 3/8$, and shows gas isodensities coincident with the dark matter isopotentials.

5.6 Asymmetric Initial Conditions: WVT Logo

Figure 8 shows another example for an arbitrarily complex setup. The top panel shows the letters “WVT” used in 3 dimensions, the lower panel gives the same example in 2 dimensions. Note that we omitted the largest SPH particles (white particles in the 2D version) in the 3D version for clarity. WVT has no difficulties in matching the desired resolution even in these complex test cases.

5.7 Mixing in Supernovae

As the shock moves out through a star in a supernova explosion, Richtmyer-Meshkov and Rayleigh-Taylor instabilities drive turbulence behind the shock. This turbulence mixes elements, dredging up radioactive ^{56}Ni and injecting hydrogen into slower moving layers. This mixing is observed in supernova light-curves and in the knots in supernova remnants. But modeling this turbulence is not trivial; both cartesian grid Eulerian and SPH codes introduce numerical turbulence (based on noise in the initial set-up) that can artificially produce spurious turbulence.

The Sedov blast wave is an ideal test for any code modeling these supernova explosions; an analytic solution exists and can be compared to simulation results. This test also exposes consequences of choices in initial conditions beyond just total particle or mesh cell count; the shock is often unstable to hydrodynamic instabilities (e.g. Richtmyer Meshkov and Rayleigh Taylor behind the shock) and any initial density perturbation introduced by the setup (or grid effects in an Eulerian code) can artificially seed this convection.

Figure 9 shows the particle distribution, in terms of mass density versus radius, at a time

$t = 0.06317$ after the launch of the shock, for simulations using concentric shell, hexagonal close-packed, and WVT initial conditions. The black line in the figure indicates the analytic solution at this time. All three sets of initial conditions used approximately 1.5 million particles (1.52 million for the shell setup, 1.50 million for the hexagonal close-packed and WVT configurations). All three simulations used the same gamma-law equation of state, with $\gamma = 7/5$.

The hexagonal close-packed arrangement used a uniform spacing of 0.01 between closest neighbors; with that spacing, 1.50 million particles extend out to a radius $r = 0.63$. Particles in the lattice had an average of 54 neighbors. Both the concentric shell arrangement and the WVT setup allow particle size to vary with radius, so that resolution can be focused where it will be most useful; consequently, they both extend over a larger range of radii and place more particles near the origin. The concentric shell arrangement extended inward to $r_{min} = 1.70 \times 10^{-3}$, where the smallest particles had smoothing lengths $h_{min} = 7.81 \times 10^{-3}$, while the WVT arrangement extended to $r_{min} = 2.56 \times 10^{-4}$, where $h_{min} = 2.69 \times 10^{-4}$. Both configurations extended outward to $r_{max} = 1$, where the largest particles had smoothing lengths $h_{max} = 2.62 \times 10^{-2}$ in the concentric shell setup and $h_{max} = 3.29 \times 10^{-2}$ in the WVT setup. Particles in the concentric shell setup had an average of 76 neighbors, while particles in the WVT setup had an average of 50 neighbors.

In the same way that initiating a Sedov blast wave calculation by injecting energy into a single mesh cell could imprint the mesh geometry onto the resulting shock, initiating a calculation by injecting energy into a single SPH particle at the origin could produce an aspherical explosion imprinted with artifacts of the arrangement of neighboring particles. In each of the three simulations, energy $E = 1$ was injected into a small spherical volume at the center of the simulation at $t = 0$. The radius of that volume varied among the simulations, according to the competing constraints that it be as small as possible, to initiate a point-like explosion, but large enough to extend out to several times the smoothing length of the innermost particles, to eliminate relics of the specific particle arrangement around the origin.

In the hexagonal close-packed configuration, simulating a point-like explosion with limited resolution near the origin became a primary concern; energy was smoothed over particles at radii $r < 0.024$, which included 81 particles. In the concentric shell setup, energy was smoothed over particles at radii $r < 3.9 \times 10^{-3}$, or the innermost 10 shells; in the WVT setup, energy was smoothed over particles at radii $r < 5.0 \times 10^{-3}$, which included 38, 148 particles.

As the shock expands outward, the variation of resolution with radius among the three sets of initial conditions becomes apparent; at $t = 0.06317$, there are 481, 562 shocked particles in the

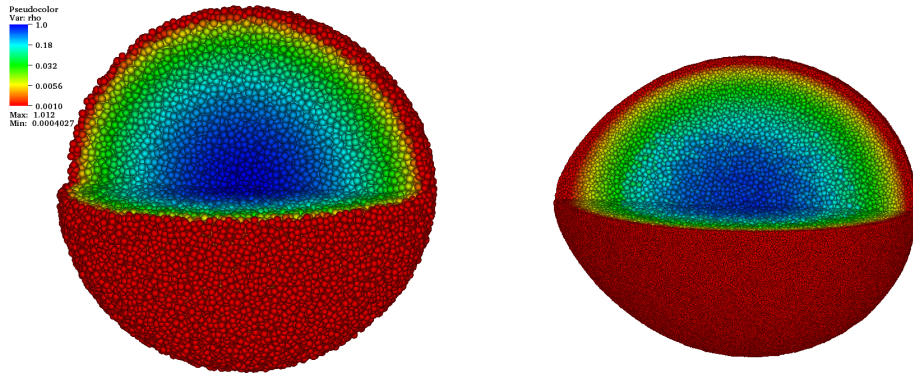


Figure 6. This three-dimensional example shows an asymmetric WVT setup for a double degenerate merger simulation. In this example, the accretor (left) is modeled with a constant particle density, whereas the donor (right) has significant more resolution in the outer layers than in the center, making SPH simulation of Roche lobe overflow feasible.

simulation using the shell setup, only 190,472 in the hex close-packed simulation, and 977,512 in the WVT simulation.

The nature of the concentric shell setup introduces a density perturbation within each shell. This perturbation grows when the shock passes through it, driving strong density perturbations and convection. With this setup, scientists were forced to drive up the resolution to ensure the initial perturbation is less than the observed (Fryer et al. 2007). For hexagonal close-packing schemes, the low resolution in the energy-injection region leads to a range of shock velocities driving strong density perturbations. WVT produces much less scatter with the same resolution, limiting the numerically-seeded turbulence in this problem.

6 COMPARISON TESTS

6.1 Interpolation Accuracy: Uniform Density

An important performance test for any SPH setup method is to find out how well it reproduces a given density field. This test will reveal the amount of perturbations that are introduced by the setup, which could seed convection, excite sound waves or trigger instabilities. The simplest such test is to see how well each method can mimic a uniform density field with a uniform particle distribution. Thus, each particle should have the same mass, which we will assume to be 1, and the same smoothing length/resolution. At the same time, this test will then provide a means to test the accuracy of the particle density distribution itself.

Figure 10 shows the the accuracy of all uniform density methods that are mentioned in the review section 3, along with the new WVT method. Each panel shows an interpolation of a unit cube containing 8,000 particles onto the x–y plane according to the standard spline SPH kernel

[htbp]

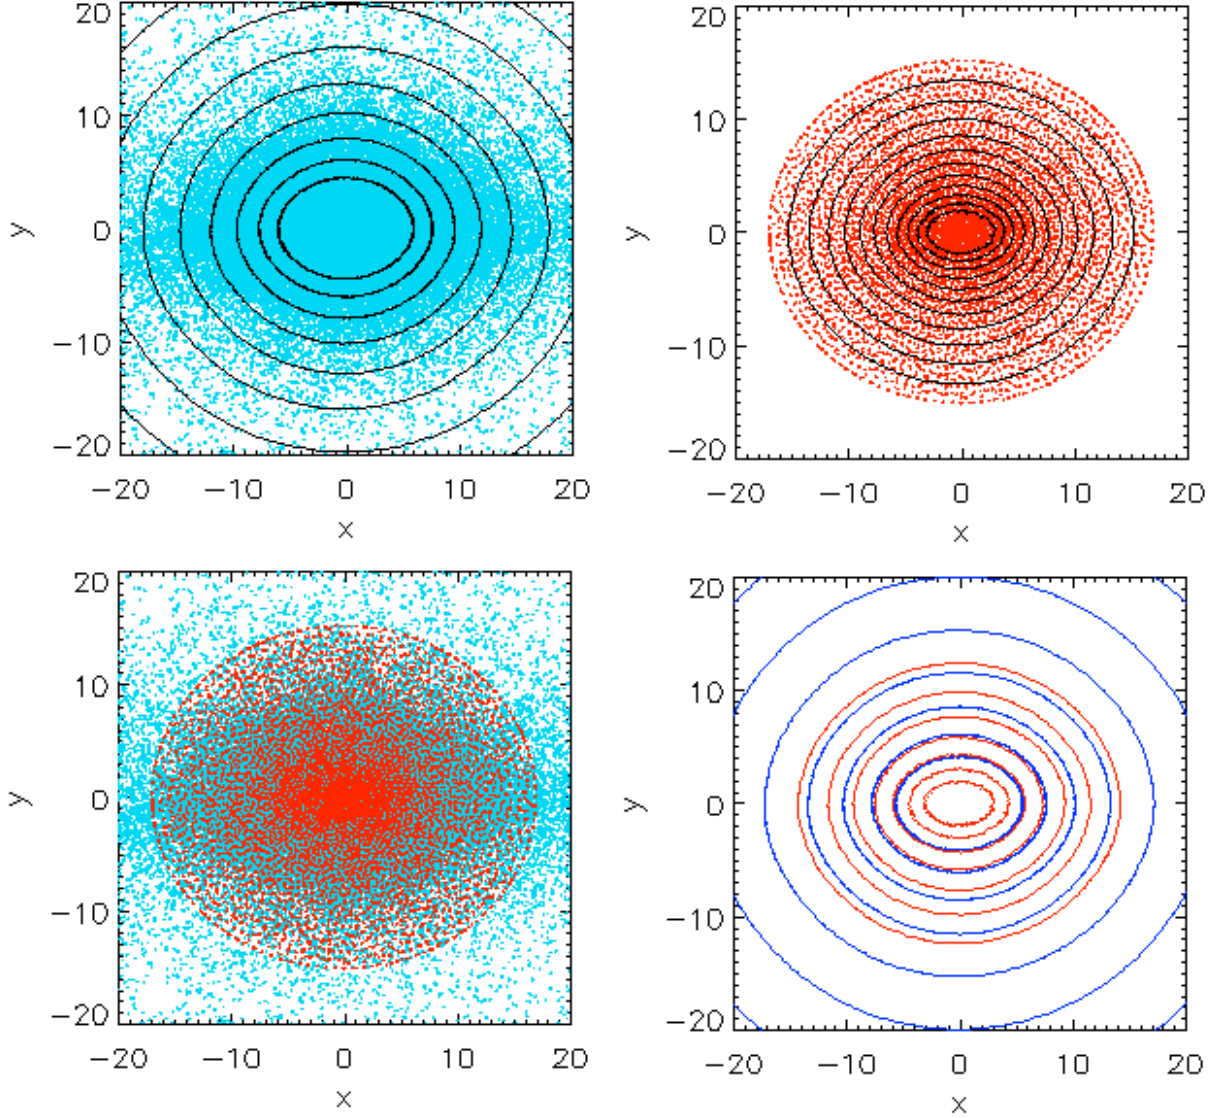


Figure 7. WVT results for a gas cloud with polytropic index of $n = 3/8$, embedded within a triaxial (axis ratio 17:15:14) dark matter potential, a slice through the simulation at $z = 0$. Top-left: Dark matter particles in cyan, with surfaces of constant potential overplotted. Top-right: SPH gas particles in red, with surfaces of constant mass density overplotted. Bottom-left: SPH gas particles in red and dark matter particles in cyan. Bottom-right: The surfaces of constant dark matter potential (blue) coincide with the surfaces of constant gas mass density (red).

targeted at containing approximately 128 neighbors. We divided each figure into two parts, with the colors in the left half showing up to 5% deviations (negative: blue, positive: red, accurate: green), whereas the right half reveals lower level (up to 1%) deviations.

The first 3 panels show the lattice configuration (cubic lattice, cubic close packing, and hexagonal close packing) which obviously have excellent interpolation properties. This is not surprising, as they are designed to be as uniform as possible, and each particle has an identical amount of neighbors. Only in the right panels, low-level noise becomes visible, in fact revealing the underlying lattice structures. Note that the cubic close packing panel shows hexagonal structures, as the

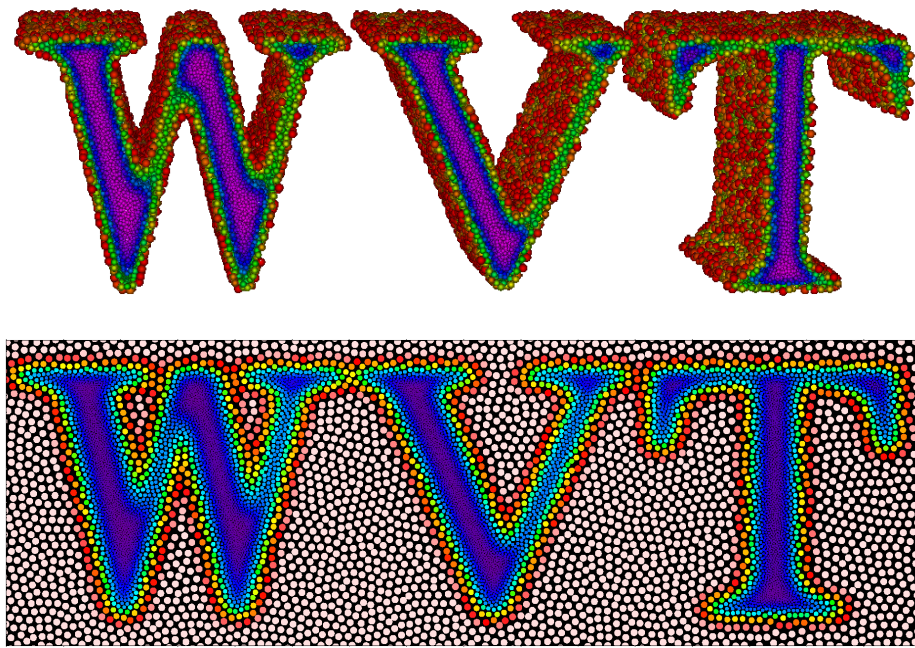


Figure 8. Examples for an arbitrary spatial configuration. The top panel shows a 3d configuration, the lower panel shows a two-dimensional configuration. Particles with large smoothing lengths (shown in white in bottom panel) are omitted in the three-dimensional view for clarity. This particular configuration shows a dynamic range of ~ 10 . Smoothing lengths are indicated by color and proportional to their symbols' sizes.

x–y plane cuts through a hexagonal layer. Other orientations of the plane would reveal different patterns, but give the same qualitative impression.

The fourth panel in the top row, shows the quaquaversal tiling configuration, which shows very strong clustering of particles, and thus deviations from the ideal values. Even with 128 neighbors, density fluctuations on the order of 5% are found throughout the simulation volume. In addition, these deviations are strongly spatially correlated, which makes numerical artifacts likely. This reason alone is grounds enough not to use quaquaversal tiling for SPH setups. This effect has also recently been pointed out by Wang & White (2007), who found that the quaquaversal tiling setup in cosmological situations leads to a spurious amount of small halos and clumping. In fact, the only setup method that produces stronger density fluctuations is by placing all particles randomly (top right panel) throughout the simulation domain.

The first panel in the lower row on the other hand shows the interpolation accuracy of a much better setup method, the shell setup. Note that the center in the shell setup method has to be treated specially to produce a uniform density for the inner shell. To achieve a uniform distribution, the inner shell would need as little as 16 particles distributed uniformly across the sphere, at which point the premise of considering this configuration to be a shell is no longer appropriate. Otherwise, with more particles on the innermost shell, the setup is bound to have a hole in the center. While in principle this hole could be filled by placing a lattice or glass structure inside

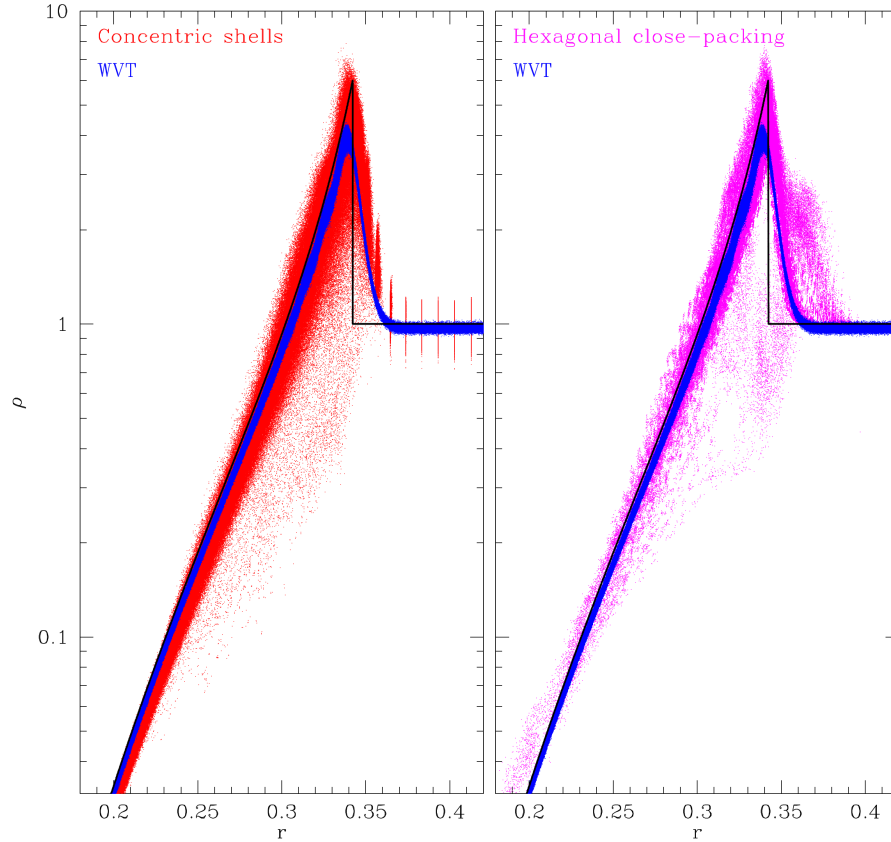


Figure 9. Density versus radius of a Sedov blast wave problem comparing the results from a 1.5 million particle WVT setup with one using concentric shells (left) and hexagonal close-packing (right). The black line indicates the analytic solution. The initial density perturbations in the concentric shell setup grow in the shock producing a broad range of particle densities. The low-resolution at the energy source leads to velocity perturbations that then create density perturbations.

the inner shell, this would likely produce artifacts at the boundaries and is thus undesirable as well. However, this method has so far never been used for such a situation, but is usually input for supernovae simulation, which have a special boundary condition at the center. Thus, we restrict our interpolation analysis in this case to an off-center location, with the center being located to the left of the shown image. Note how you can pick out individual shells in the noise. The level of noise in the density interpolation for 128 neighbors is on the order of 1%, consistent with findings by Fryer et al. (2007). However, note that the setup is only optimized within one shell, which leads to the noise deviations having a preferred radial direction perpendicular to the shell structures.

Within one shell, the setup has similar properties to a uniform gravitational glass, shown in the second panel of the bottom row. The level of noise for this method is very low, generally on the order of 0.5% at most. Also note that the noise is isotropic with no preferred direction, as is true for

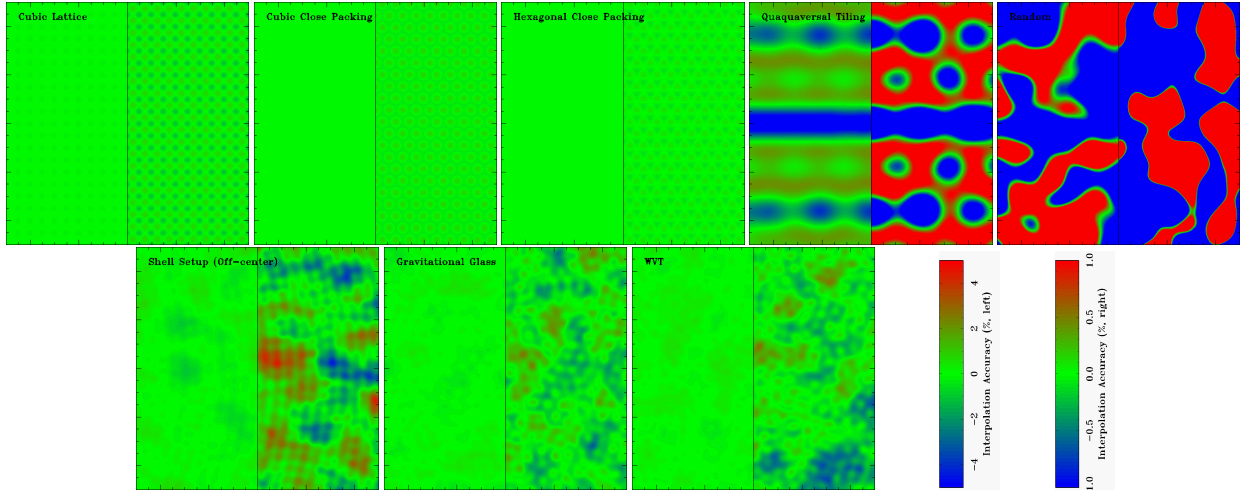


Figure 10. Comparison of the interpolation accuracy for 128 neighbors in the cubic lattice, cubic close packing, hexagonal close packing, quaquaversal tiling, random configuration, shell setup, gravitational glass, and the new WVT approach (top left to bottom right). Colors indicate deviations from the target density, with blue colors showing negative and red colors denoting positive deviations. Each panel is divided into two halves with different dynamic ranges: $\pm 5\%$ on the left and $\pm 1\%$ on the right side. Note that the shell setup is shown at an off-center location, to avoid the discussed special treatment of the center in this comparison. Quaquaversal tiling and the random setup perform noticeably worse than any other method, whereas the uniform grid setups perform best as expected. The non-gridded setup methods perform equally well, with very high interpolation accuracies that never exceed 1%.

the underlying particle distribution. These desirable properties make the gravitational glass setup a suitable choice for uniform density distributions.

The last panel in the bottom row shows the interpolation properties of our new WVT setup method. Note the similarity to the gravitational glass, with an equally low amount of noise and an isotropic distribution of the noise without preferred directions. In our WVT implementation, each repulsive force is roughly proportional to r^{-2} , and the ratio of scale lengths is 1 for a uniform distribution, which makes the WVT setup locally very similar to a gravitational glass and its characteristics.

Let us quickly summarize our comparison of the interpolation accuracy of the various methods. The three uniform lattices (cubic lattice, cubic close packing, hexagonal close packing) have excellent interpolation characteristics for a uniform density and should be used in situations where lattice effects are expected to be unimportant, but a very low level of initial numerical noise is needed. The quaquaversal and random initial condition are unacceptable for any application due to their low level of interpolation accuracy. In spherical symmetry, the shell setup provides an adequate configuration, but without a special treatment of the center is unable to reproduce a solid, uniform center of the sphere. The gravitational glass and the new WVT setup method both performed best in the interpolation test for non-lattice configurations. Both have a high level of interpolation accuracy with maximum deviations generally on the order of 0.5% for 128 neighbors.

6.2 Interpolation Accuracy: Non-Uniform Density

We now consider a very similar test, but for a non-uniform particle distribution. We choose to test all adaptive setup methods listed in section 3 along with our new WVT setup. As a comparison test, we choose a spherical setup with the resolution intended to scale as $r^{2/3}$. Figure 11 shows the comparison data, in a similar fashion to Figure 10.

Even with this moderate amount of stretching, all the lattice configurations (stretched cubic lattice, stretched cubic close packing, stretched hexagonal close packing) perform very poorly in this test. This is not surprising when one considers the 3d structure of the stretched lattices in Figure 2. The problem for these structures is that the stretching factor is a function of radius, and thus not parallel to one of the lattice axes. Thus, originally parallel planes get strongly bent during the stretch process, and the particle spacings within one such multiplied by different stretching factors. This results in a very uneven distribution of particle density, and the lattice structure can now be easily picked out in the first three panels of Figure 11.

The quaquaversal tiling and random configuration perform even more poorly, and the problems of the uniform density test get even more exaggerated.

The intrinsically adaptive shell setup (bottom row, left) performs very well in this test, with density perturbations equivalent to the uniform density test, usually not exceeding 1% for 128 neighbors. The only disadvantage of the shell setup is that the density deviations are systematically in the radial direction, as it is only optimized within a shell.

Interestingly and maybe surprisingly, the stretched glass on the other hand performs rather poorly, as shown in the middle panel of the bottom row. Even though the gravitational glass has excellent interpolation properties for uniform densities, this is not the case when stretched in a radial (or any other) direction. As we had already noticed in the lattice configurations, the stretching procedure tends to pronounce voids between planes that are perpendicular to the direction in which the stretching is applied. The glass does not have a uniform clear lattice structure, but tends to order particles along randomly oriented strings on a local level, as can be seen in the left panel of Figure 12. Thus the stretching will preferably pick out those features that are perpendicular to our stretching direction, i.e. those inside shells. This leads to the wavy shell-like features in the right panel of Figure 12 which shows the stretched glass. If you look closely, you can pick out the strings in the left panel, which survive and produce the shells on the right. In this particular example, the density deviations are on the order of 3%, which is significantly less than the lattice structures. However, with stronger stretching, these features will only become even more pronounced.

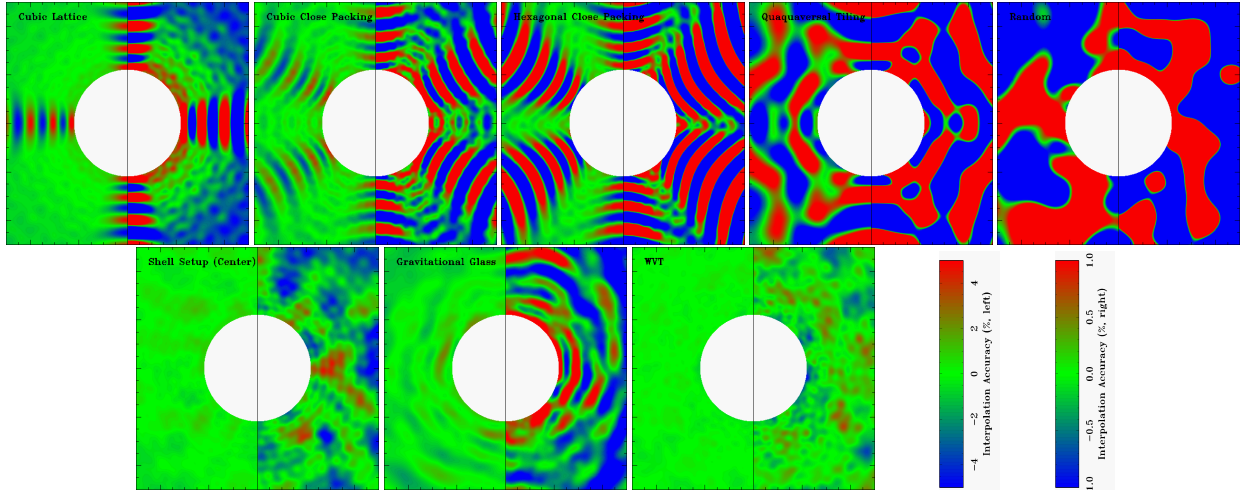


Figure 11. The equivalent of Figure 10, but for spatially adaptive setups: stretched cubic lattice, stretched cubic close packing, stretched hexagonal close packing, stretched quaquaversal tiling, random configuration, shell setup, stretched gravitational glass, and the new WVT approach (top left to bottom right). All uniform-grid-based setups, quaquaversal tiling and the random configuration perform very badly. Of the non-gridded setup methods, the WVT setup performs best with density inaccuracies below 1%. The stretched gravitational glass introduces artifacts that are located inside shells, whereas the shell setup preferably has deviations in the radial direction.

The WVT setup (bottom row, far right panel) does not exhibit these features. Note how the level of noise for the adaptive setup is as low as for the uniform distribution. This is due to the fact that the WVT setup knows beforehand the desired resolution at each point in space, and is not a stretched version of a uniform distribution, allowing to converge to the optimal solution in either case. We also note that this high level of interpolation accuracy does not depend on the spherical symmetry as in the shell setup, which makes WVT much more versatile.

In summary, we find that stretched lattice configurations are not suitable for producing non-uniform particle distributions. Quaquaversal tiling and random configuration have even worse interpolation properties. The shell setup has acceptable levels of noise, but is restricted to spherical symmetry and the noise distribution is preferably along the radial direction. Surprisingly, the stretched glass setup also introduces artifacts but rather in shells which could lead to radial pulsations in an SPH simulation. The WVT setup on the other hand has interpolation properties that are equivalent from the uniform distribution, which makes it the method of choice for adaptive resolution requirements. All other distributions should be relaxed into their equilibrium configuration prior to being used.

6.3 Particle Noise

Another way of judging the characteristics of an SPH setup is to measure the particle noise inside a uniform distribution of particles within a uniform density distribution. In an ideal situation, all the pressure forces of individual particles cancel out, and the net force is 0 or very small compared

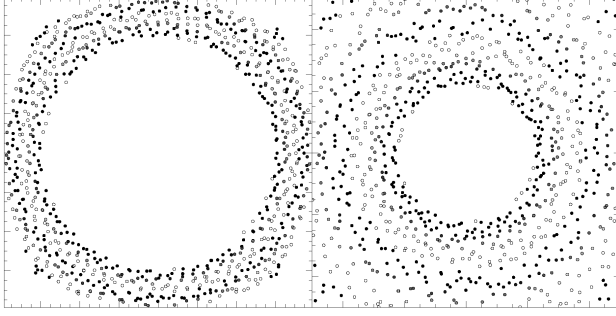


Figure 12. Unstretched (left) and stretched glass (right). Note the wavy shell-like structure in the stretched glass. To bring out this structure, particles are periodically colored by their radius in the unstretched glass, which is then also applied to the stretched glass.

to an individual force. Then, the contribution to the pressure force of an individual SPH particle i on another particle j is (Monaghan 1992)

$$\frac{dv_{ij}}{dt} \propto -m_i \frac{P_i}{\rho_i^2} \nabla_i W_{ij}. \quad (4)$$

For our test, we set up conditions such that the particle masses $m_i = 1$, the average pressure $\bar{P}_i = 1$ and the average density $\bar{\rho}_i = 1$. Thus, we expect a single particle to contribute $|\nabla W_{ij}| \approx 0.08$ on average to the acceleration term. Assuming Poisson noise, we thus expect the noise for a random configuration to be on the order of $(N_{\text{neigh}})^{1/2} |\nabla W_{ij}| = 0.64$ for 64 neighbors. Our random configuration yields a value of 0.60 ± 0.25 in our test setup, consistent with expectations. We will take this measured value of 0.6 as a reference point to measure the performance of the other setup methods in terms of “fractional Poisson noise”.

As expected, the uniform lattice configurations yield a perfect equilibrium down to machine precision. The quaquaversal tiling on the other hand shows very poor performance again, with a particle noise of about 30% the Poisson level, consistent with our findings in the density interpolation accuracy. Both the shell setup ($3.9 \pm 1.7\%$), the gravitational glass ($3.1 \pm 1.3\%$) and WVT ($3.9 \pm 1.7\%$) reduce the noise by an order of magnitude. All quoted noise values are measured for 64 neighbors, and can be further reduced by increasing the number of neighbors. We expect the noise to drop as $N_{\text{neigh}}^{-1/2}$, i.e. doubling the number of neighbors will decrease the numerical noise by only about 30%.

6.4 Summary

In Table ??, we summarize our findings about positive and negative characteristics of all considered setup methods for uniform particle distributions. We also give recommendations which method should be preferred in which situation.

Table 1. Comparison of SPH Particle Setup Methods.

Method	Adaptive?	Ease of Use	Noise	Advantages	Disadvantages	Recommendation
Cubic Lattice	no	- very fast - very easy to implement	very low	- very good interpolation properties - trivial to implement - not an optimal packing scheme	- prone to artifacts due to very strong particle alignment along grid axes - applicable to uniform distributions only - unstable equilibrium configuration	- Should not be used, as it is affected by strong lattice effects and does not present a stable equilibrium configuration
Cubic Close Packing	no	- very fast - easy to implement	very low	- optimal packing scheme - stable equilibrium configuration	- applicable to uniform distributions only - prone to artifacts due to strong particle alignment along lattice symmetry axes	- Can be used for mainly static problems that do not collapse or exhibit shocks (lattice effects), and if interpolation accuracy is important
Hexagonal Close Packing	no	- very fast - easy to implement	very low	- optimal packing scheme - stable equilibrium configuration	- applicable to uniform distributions only - prone to artifacts due to strong particle alignment along lattice symmetry axes	- Can be used for mainly static problems that do not collapse or exhibit shocks (lattice effects), and if interpolation accuracy is important
Quaquaversal Tiling	no	fast & easy	high	- none	- non-intuitive to understand lattice effects - applicable to uniform distribution only	- We do not recommend to use this method for SPH
Random Configuration	yes	very fast & easy	very high	- applicable to arbitrary geometries	- very bad interpolation properties - very clumpy - dominated by Poisson noise	- We do not recommend to use this method for SPH
Gravitational Glass	no	- convergence slow - publicly available as part of Gadget2	low	- very good interpolation properties - no lattice structures - with periodic boundaries, one box can be used to generate arbitrarily large setups - noise isotropic	- applicable to uniform distribution only	- This is a very good choice for uniform particle distributions if interpolation accuracy is important and lattice effects need to be avoided
Stretched Lattice	yes	- very fast - easy to implement	very high	- fast and easy to implement	- strongly anisotropic particle distribution due to lattice effects - not a stable configuration, has to be relaxed into equilibrium configuration	- This setup has to be used in conjunction with the relaxation technique, and should never be used directly
Stretched Glass	yes	- fast (if glass available) - easy to implement	high	- better starting point for initial conditions than a stretched lattice	- noise non-isotropic, tends to be spatially correlated in shells - not a stable configuration, has to be relaxed into equilibrium configuration	- This setup has to be used in conjunction with the relaxation technique, and should never be used directly, but faster to relax than a stretched lattice
Shell Setup	yes	- very fast if shell templates are already available	low	- very well suited to produce perfect spherical symmetry - very good interpolation properties	- computationally expensive to generate shell templates - requires a lot of shell templates - only optimized within shells - noise non-isotropic	- Good choice for spherical symmetric initial conditions if interpolations accuracy is important, but center needs special treatment.
WVT Setup	yes	- fast & easy - publicly available	low	- applicable to arbitrary geometries - serial version available, parallel version easy to implement on top of an existing SPH code - very good interpolation properties	- only serial version available	- Good choice for any uniform or adaptive setup. Only choice for adaptive setups without spherical symmetry.

Figure 13. Table

7 CONCLUSIONS

We have presented an extensive comparison of all setup method currently employed in astrophysics that we are aware of. In particular, we discuss spatially uniform configurations such as a cubic lattice, cubic close packing, hexagonal close packing, quaquaversal tiling, and gravitational glasses. For spatially adaptive methods, we also include the random probability distribution, stretched lattice, stretched glass, and a concentric shell setup. To the best of our knowledge, the stretched glass and concentric shell setup has not been described in the literature before.

The main focus of our paper however is a new setup method based on weighted Voronoi tessellations. This new method allows for arbitrary geometrical configurations, which has not been possible before. We show that this new method is easy to implement on existing SPH codes and demonstrate its superior characteristics in several examples. In principle, it would be straightforward to extend this method to adaptively bin ungridded n -dimensional data.

This method has now been used in a variety of astrophysics problems from core-collapse supernovae (Ellinger et al. in preparation) to modeling binary interactions (Raskin et al. 2009, 2010; Fryer et al. 2010). Especially in doing binary interaction calculations comparing multiple techniques, using identical initial conditions is critical and these initial conditions tend to have complicated structures caused by tidal effects (Passy et al., in preparation, Motl et al., in preparation).

ACKNOWLEDGMENTS

This work was carried out in part under the auspices of the National Nuclear Security Administration of the Department of Energy (DOE) at Los Alamos National Laboratory and supported by a LANL/LDRD program. Three-dimensional images were created using the `visIt` package developed at LLNL.

REFERENCES

- Benz W., Hills J. G., 1987, *ApJ*, 323, 614
- Benz W., Slattery W. L., Cameron A. G. W., 1986, *Icarus*, 66, 515
- Bonnell I., Bastien P., 1992, *ApJ*, 401, 654
- Cappellari M., Copin Y., 2003, *MNRAS*, 342, 345
- Conway J. H., Radin C., 1998, *Inventiones Mathematicae*, 132, 179
- Cox T. J., Jonsson P., Primack J. R., Somerville R. S., 2006, *MNRAS*, 373, 1013

- Cuadra J., Nayakshin S., Springel V., Di Matteo T., 2005, MNRAS, 360, L55
- Davies M. B., Benz W., Hills J. G., 1991, ApJ, 381, 449
- Davies M. B., Benz W., Hills J. G., 1992, ApJ, 401, 246
- Davies M. B., Benz W., Piran T., Thielemann F. K., 1994, ApJ, 431, 742
- Diehl S., Fryer C. L., Herwig F., Oishi J. S., Mac Low M.-M., De Marco O., 2008, ApJ, in preparation
- Motl P. M., Tohline J. E., Fryer C. L., et al. 2011, ApJ, in preparation
- Diehl S., Statler T. S., 2006, MNRAS, 368, 497
- Efstathiou G., Davis M., White S. D. M., Frenk C. S., 1985, ApJS, 57, 241
- Ellinger C.I., Young P.A., Fryer C. L., Rockefeller G., submitted to ApJ
- Fryer C. L., Heger A., 2005, ApJ, 623, 302
- Fryer C. L., Hungerford A. L., Rockefeller G., 2007, International Journal of Modern Physics D, 16, 941
- Fryer C. L., Rockefeller G., Warren M. S., 2006, ApJ, 643, 292
- Fryer C. L., Warren M. S., 2002, ApJ, 574, L65
- Fryer C. L., Young P. A., 2007, ApJ, 659, 1438
- Fryer, C.L., et al., 2010, ApJ, 725, 296
- Gingold R. A., Monaghan J. J., 1977, MNRAS, 181, 375
- Hansen S. H., Agertz O., Joyce M., Stadel J., Moore B., Potter D., 2007, ApJ, 656, 631
- Herant M., 1994, Memorie della Societa Astronomica Italiana, 65, 1013
- Hungerford A. L., Fryer C. L., Warren M. S., 2003, ApJ, 594, 390
- Holley-Bockelmann, K., Mihos, J. C., Sigurdsson, S., & Hernquist, L. 2001, ApJ, 549, 862
- Kazantzidis, S., Magorrian, J., & Moore, B. 2004, ApJ, 601, 37
- Khalatyan A., Cattaneo A., Schramm M., Gottlöber S., Steinmetz M., Wisotzki L., 2008, MNRAS, 387, 13
- Lee W. H., Kluzniak W., Nix J., 2001, Acta Astronomica, 51, 331
- Lloyd S., 1982, IEEE Trans. Inform. Theory, 28, 129
- Lombardi J. C., Sills A., Rasio F. A., Shapiro S. L., 1999, Journal of Computational Physics, 152, 687
- Lucy L. B., 1977, AJ, 82, 1013
- Mayer L., Quinn T., Wadsley J., Stadel J., 2002, Science, 298, 1756
- Merritt, D. 1985a, MNRAS, 214, 25P
- . 1985b, AJ, 90, 1027

- Møller J., 1994, *Lectures on Random Voronoi Tessellations*. Springer-Verlag, New York
- Monaghan J. J., 1992, *ARA&A*, 30, 543
- Monaghan J. J., Lattanzio J. C., 1991, *ApJ*, 375, 177
- Navarro J. F., Frenk C. S., White S. D. M., 1995a, *MNRAS*, 275, 720
- Navarro J. F., Frenk C. S., White S. D. M., 1995b, *MNRAS*, 275, 56
- Nelson A. F., Benz W., 2003, *ApJ*, 589, 556
- Osipkov, L. P. 1979, *Soviet Astronomy Letters*, 5, 42
- Passy J.-C., De Marco O., Fryer C. L., Herwig, F., Diehl, S., Oishi, J.S., Mac Low M.-M., Bryan G.L., Rockefeller G., accepted by *ApJ*, astro-ph/1107.5072
- Rasio F. A., Shapiro S. L., 1991, *ApJ*, 377, 559
- Rasio F. A., Shapiro S. L., 1992, *ApJ*, 401, 226
- Raskin C., Timmes F.X., Scannapieco E., Diehl S., Fryer C. L., 2009, *ApJ*, 399, L156
- Raskin C., Timmes F.X., Scannapieco E., Diehl S., Fryer C. L., 2010, *ApJ*, 724, 111
- Rockefeller G., Fryer C. L., Baganoff F. K., Melia F., 2005, *ApJ*, 635, L141
- Rockefeller G., Fryer C. L., Melia F., Warren M. S., 2004, *ApJ*, 604, 662
- Rosswog S., Liebendörfer M., Thielemann F.-K., Davies M. B., Benz W., Piran T., 1999, *A&A*, 341, 499
- Rosswog S., Ramirez-Ruiz E., Hix R., 2008, *ArXiv e-prints*, 808
- Springel V., 2005, *MNRAS*, 364, 1105
- Terman J. L., Taam R. E., Hernquist L., 1994, *ApJ*, 422, 729
- Thakar A. R., Ryden B. S., 1998, *ApJ*, 506, 93
- Wang J., White S. D. M., 2007, *MNRAS*, 380, 93
- Whitworth A. P., 1998, *MNRAS*, 296, 442
- Widrow, L. M. 2008, *ApJ*, 679, 1232
- Yoon S.-C., Podsiadlowski P., Rosswog S., 2007, *MNRAS*, 380, 933